

**In the claims:**

Please amend the claims as follows.

1 – 43. Cancelled.

44. (Currently Amended) A method of maintaining order of transactions in a distributed communication system, the distributed communication system comprising a plurality of nodes interconnected by a plurality of communication links, the plurality of nodes having access to a plurality of addressable memory locations, the plurality of nodes comprising a source node and a target node, the method comprising the acts of:

dispatching, by the source node, a first request directed to a first memory address accessible by the target node;

transmitting, from the target node, a first response directed to the source node in response to the first request;

transmitting, from the source node, a second response directed to the target node after receipt of the first response; and

stalling service, by the target node, of a second request directed to the first memory address pending receipt of the second response; and

issuing, by the target node, a probe in response to the first request, the probe being directed to each of the plurality of nodes to determine whether any of the plurality of nodes is caching data corresponding with the first memory address; and

issuing, by each of the plurality of nodes, a third response in response to the probe,

wherein act of transmitting the second response, from the source node, is performed after receipt of all of the third responses; and

wherein each of the third responses is directed to the target node, and wherein the act of transmitting the first response directed to the source node is performed after receipt of all of the third responses by the target node.

45. (Previously Presented) The method as recited in claim 44, wherein the first request is associated with a first transaction, and the act of transmitting the first response is performed after the first transaction has reached a memory commit point.

46. (Previously Presented) The method as recited in claim 45, wherein the act of transmitting the second response is performed after the first transaction has reached a processor commit point.

47. (Previously Presented) The method as recited in claim 45 wherein the first transaction is a write transaction, wherein the first response is a Target Done response, and wherein the second response is a Source Done response.

48. (Previously Presented) The method as recited in claim 44, wherein the first request comprises a read request, and wherein the first response comprises a read response.

49. (Previously Presented) The method as recited in claim 44, wherein the first request is associated with a first transaction, and the method comprises the act of: stalling, by the source node, dispatch of a second transaction pending receipt of the first response.

50. (Cancelled)

51. (Cancelled)

52. (Currently Amended) The method as recited in claim ~~50~~ 44, wherein each of the third responses is directed to the source node.

53. (Currently Amended) The method as recited in claim ~~50~~ 44, wherein one of the third responses is a read response, the read response indicating that the node which issued the read response is storing data corresponding to the first memory address.

54. (Cancelled)

55. (Cancelled)

56. (Cancelled)

57. (Cancelled)

58. (Previously Presented) The method as recited in claim 44, wherein the source node comprises a processor.

59. (Previously Presented) The method as recited in claim 44, wherein the source node comprises a host bridge.

60. (Previously Presented) The method as recited in claim 44, wherein the target node comprises a memory controller configured to access the first memory address.

61. (Currently Amended) A method of maintaining order of transactions issued in a distributed communication system, the distributed communication system comprising a plurality of nodes interconnected by a plurality of communication links, the plurality of nodes configured to access a plurality of addressable memory locations for storing data, the plurality of nodes comprising a source node and a target node, the method comprising the acts of:

storing at the target node a first request received from the source node and directed to a first memory address;

storing at the target node a second request directed to the first memory address;  
 servicing the first request;  
 transmitting a communication to the source node in response to the act of  
     servicing the first request; and  
 stalling an act of servicing the second request pending receipt by the target node  
     of a source response transmitted from the source node in response to the  
     communication;

wherein the acts of storing the first request and the second request at the target  
 node comprises the act of ordering the first request and the second request  
 in a queue in the order in which the first request and the second request  
 were received;

wherein the first request comprises a read request, and wherein the  
 communication transmitted to the source node comprises a plurality of  
 responses issued from the plurality of nodes;

wherein the plurality of responses comprises a memory access response issued  
 from the target node, and wherein the act of servicing the first request  
 comprises the acts of:

accessing, by the target node, a memory location associated with the first  
 memory address;

transmitting to the source node the memory access response based on  
 the act of accessing by the target node; and

determining whether data corresponding to the first memory address is  
 cached at any of the plurality of nodes;

wherein the plurality of responses comprises a plurality of cache responses issued  
 by the plurality of nodes, and wherein the act of determining whether data  
 corresponding to the first memory address is cached at any of the plurality  
 of nodes comprises the acts of:

issuing a probe directed to each of the plurality of nodes; and

issuing, by each of the plurality of nodes, one of the cache responses in response to the probe, each of the cache responses being directed to the source node; and

wherein if a first node of the plurality nodes is caching data corresponding to the first memory address, the method comprises the acts of:  
issuing, by the first node, a memory cancel response directed to the target node; and  
formatting the cache response issued from the first node to indicate the act of issuing the memory cancel response.

62. (Cancelled)

63. (Previously Presented) The method as recited in claim 61, wherein the first request comprises a write request, and wherein the act of transmitting the communication to the source node is performed when the act of servicing the write request has reached a memory commit point.

64. (Previously Presented) The method as recited in claim 63, wherein the communication comprises a Target Done response, and wherein the method comprises the act of generating the Target Done response by the target node.

65. (Previously Presented) The method as recited in claim 63, wherein the act of servicing the first request comprises the act of determining whether data corresponding to the first memory address is cached at any of the plurality of nodes.

66. (Previously Presented) The method as recited in claim 65, wherein the act of serving the first request has reached the memory commit point when the act of determining whether data corresponding to the first memory address is cached at any of the plurality of nodes is complete.

67. (Previously Presented) The method as recited in claim 65, wherein the act of determining whether data corresponding to the first memory address is cached at any of the plurality of nodes comprises the acts of:

issuing a probe directed to each of the plurality of nodes; and  
issuing, by each of the plurality of nodes, a cache response in response to the probe, each cache response being directed to the target node and indicating whether data corresponding to the first memory address is cached at the particular node.

68. (Previously Presented) The method as recited in claim 67, wherein the memory commit point is reached when all of the cache responses have been received by the target node.

69. (Cancelled)

70. (Cancelled)

71. (Cancelled)

72. (Cancelled)

73. (Currently Amended) The method as recited in claim ~~72~~ 61, comprising the act of canceling the act of accessing the memory location in response to the memory cancel response.

74. (Currently Amended) The method as recited in claim ~~72~~ 61, comprising the act of:

formatting the source response to indicate the act of issuing the memory cancel response; and

wherein the act of stalling service of the second request is stalled pending receipt of the memory cancel response by the target node.

75. (Currently Amended) A communication node for a distributed communication system comprising a plurality of communication nodes interconnected by a plurality of communication links, the node comprising:

a memory controller to control access to a memory, the memory comprising a plurality of memory locations corresponding to a plurality of memory addresses;

an interface configured to connect to a communication link; and

communication logic coupled to the memory controller, and the interface, wherein the communication logic is configured to:

store a first request received from a source via the interface, the first communication being directed to a first memory address of the plurality of memory addresses;

store a second communication directed to the first memory address;

generate a first response directed to the source in response to the first request; and

stall the second request pending receipt from the source of a second response in response to the first response;

wherein the communication logic is configured to generate a probe for transmission to each of the plurality of communication nodes in the distributed communication system, the probe to determine whether data corresponding to the first memory address is cached at any of the plurality of communication nodes;

wherein the first request comprises a read request, and wherein, the communication logic is configured to issue the read request to the memory controller to access the memory location corresponding to the first memory address;

wherein the communication logic is configured to cancel the access by the memory controller to the memory location in response to a memory cancel response received from a particular communication node of the plurality of communication nodes, the memory cancel response indicating that the particular communication node is caching data corresponding to the first memory address.

76. (Previously Presented) The communication node as recited in claim 75, comprising:

a processor; and

a cache to store data, the cache being coupled to the processor and the communication logic; and

wherein the communication logic comprises a buffer configured to store the first request and the second request in the order received.

77. (Cancelled)

78. (Currently Amended) The communication node as recited in claim ~~77~~ 75, wherein the first request comprises a write request, and wherein the communication logic is configured to receive a plurality of cache responses in response to the probe, each cache response indicating whether data corresponding to the first memory address is cached at a particular communication node of the plurality of communication nodes in the distributed communication system.

79. (Previously Presented) The communication node as recited in claim 78, wherein the communication logic is configured to generate the first response directed to the source when all of the plurality of cache responses to the probe have been received.

80. (Previously Presented) The communication node as recited in claim 79, wherein the first response comprises a Target Done response.



81. (Cancelled).

82. (Currently Amended) The communication node as recited in claim ~~81~~ 75, wherein the first response comprises a read response in response to the access to the memory location.

83. (Cancelled)

84. (Previously Presented) The communication node as recited in claim 83, wherein the communication logic is configured to:  
determine, based on the second response received from the source, whether the memory cancel response was issued; and  
stall the second request pending receipt of the memory cancel response.

85. (Previously Presented) The communication node as recited in claim 75, wherein the second response comprises a Source Done response.

86. (Previously Presented) The communication node as recited in claim 75, wherein the communication logic comprises packet-based communication logic.

87. (Currently Amended) A distributed communication system, comprising:  
a plurality of nodes, the plurality of nodes comprising a source node and a target node;  
a plurality of communication links interconnecting the plurality of nodes; and  
a memory accessible by the plurality of nodes, the memory comprising a plurality of memory locations corresponding to a plurality of memory addresses, wherein each of the plurality of nodes is configured to control access to a portion of the memory locations;

wherein the source node is configured to dispatch a first request directed to a first memory address accessible by the target node;

wherein the target node is configured to transmit a first response to the source node in response to the first request;

wherein the source node is configured to transmit a second response to the target node in response to the first response; and

wherein the target node is configured to stall service of a second request directed to the first memory address pending receipt of the second response;

wherein each of the plurality of nodes comprises a cache to store data, and wherein the target node is configured to issue a probe in response to the first request, the target node directing the probe to each of the plurality of nodes to determine whether data corresponding to the first memory address is stored in the cache of any of the plurality of nodes, and wherein each of the plurality of nodes is configured to issue a cache response in response to the probe, the cache response indicating whether data is stored in the cache of the respective node;

wherein the first request comprises a read request, and wherein a first node of the plurality of nodes is configured to issue a memory cancel response directed to the target node if the first node is storing data corresponding to the first memory address in its cache;

wherein the first node is configured to format the respective cache response to indicate that the first node has issued the memory cancel response;

wherein the source node is configured to format the second response to indicate issuance of the memory cancel response based on the cache response; and

wherein, in response to the second response, the target node is configured to stall the second request pending receipt of the memory cancel response.

88. (Previously Presented) The system as recited in claim 87, wherein the first request is associated with a first transaction, and wherein the source node is configured to stall dispatch of a second transaction pending receipt of the first response.

89. (Previously Presented) The system as recited in claim 87, wherein the first request comprises a write request, and wherein the target node is configured to transmit the first response when the first request reaches a memory commit point.

90. (Cancelled.)

91. (Currently Amended) The system as recited in claim ~~90~~87, wherein the target node is configured to format the probe, based on the first request, such that the probe identifies a destination for the plurality of cache response.

92. (Previously Presented) The system as recited in claim 91, wherein the target node formats the probe to identify the target node as the destination if the first request comprises a write request.

93. (Previously Presented) The system as recited in claim 91, wherein the target node formats the probe to identify the source node as the destination if the first request comprises a read request.

94. (Previously Presented) The system as recited in claim 93, wherein if the first request comprises a read request, the target node issues the first response directed to the source node after all of the cache responses are received by the target node.

95. (Previously Presented) The system as recited in claim 92, wherein if the first request comprises a write request, the source node issues the second response

directed to the target node after all of the caches responses and the first response are received by the source node.

96. (Cancelled)

97. (Cancelled)

98. (Currently Amended) The system as recited in claim 96 87, wherein the target node is configured to cancel the access to the memory location corresponding to the first memory address in response to the memory cancel response.

99. (Previously Presented) The system as recited in claim 87, wherein the source node comprises a host bridge.

100. (New) A method of maintaining order of transactions in a distributed communication system, the distributed communication system comprising a plurality of nodes interconnected by a plurality of communication links, the plurality of nodes having access to a plurality of addressable memory locations, the plurality of nodes comprising a source node and a target node, the method comprising the acts of:

dispatching, by the source node, a first request directed to a first memory address accessible by the target node;

transmitting, from the target node, a first response directed to the source node in response to the first request;

transmitting, from the source node, a second response directed to the target node after receipt of the first response; and

stalling service, by the target node, of a second request directed to the first memory address pending receipt of the second response;

issuing, by the target node, a probe in response to the first request, the probe being directed to each of the plurality of nodes to determine whether any of the plurality of nodes is caching data corresponding with the first memory address; and  
issuing, by each of the plurality of nodes, a third response in response to the probe,  
wherein act of transmitting the second response, from the source node, is performed after receipt of all of the third responses;  
wherein the first request comprises a read request, and wherein if the probe determines that a first node of the plurality of nodes is caching data corresponding to the first address, then the method comprises the act of:  
issuing, by the first node, a memory cancel response directed to the target node to cancel a memory access by the target node to the first memory address.

101. (New) The method as recited in claim 100, comprising the acts of:  
canceling the memory access; and  
issuing by the target node a target done response directed to the source node in response to the memory cancel response.

102. (New) The method as recited in claim 100, wherein the act of issuing the third response by the first node in response to the probe comprises the acts of:  
formatting the third response to indicate the issuance of the memory cancel response by the first node; and  
formatting the second response to indicate the issuance of the memory cancel response.

103. (New) The method as recited in claim 100, wherein the act of stalling service of the second request comprises stalling service pending receipt of the memory cancel response by the target node.